Cell Reports
Physical Science

CellPress
OPEN ACCESS

Article

Generation of dual-target compounds using a transformer chemical language model

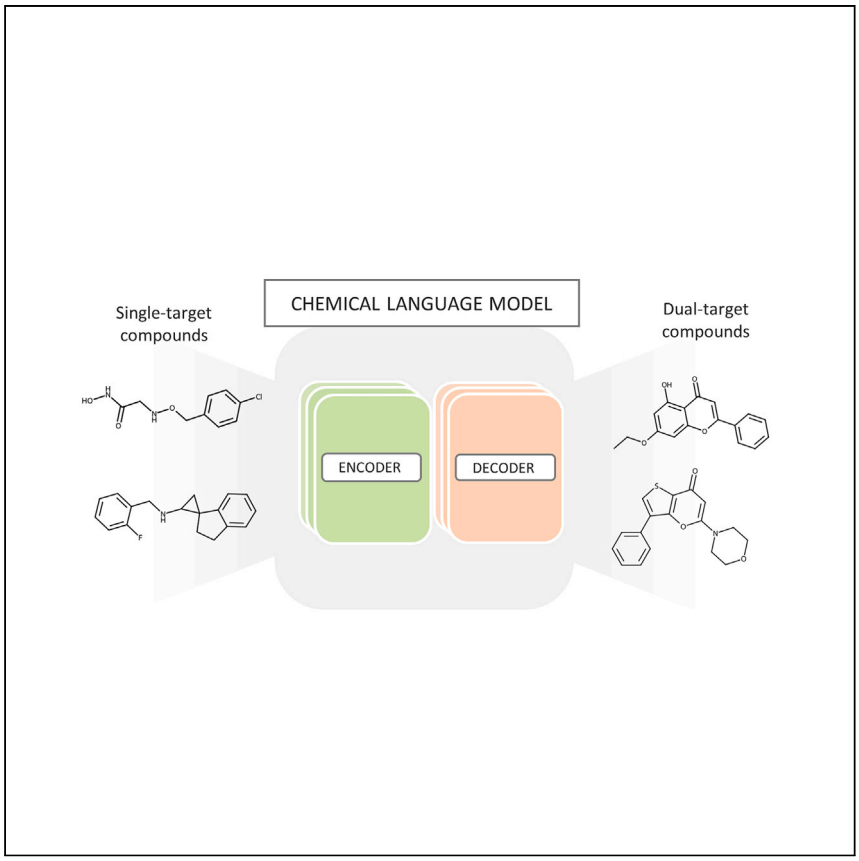

Srinivasan and Bajorath report transformer-based chemical language models for the generation of compounds with defined activity against unrelated targets. Using a new cross-fine-tuning technique, model variants are optimized for the design task by applying molecular similarity constraints; fine-tuned models chart relevant chemical space and exactly reproduce known dual-target compounds.

Sanjana Srinivasan, Jürgen Bajorath

bajorath@bit.uni-bonn.de

Highlights

Transformer models are derived for generating compounds with dual-target activity

Different pre-trained and fine-tuned model versions are evaluated

Known dual-target compounds not encountered during training are exactly reproduced

Similarity of single- and dual-target compounds is critical for model derivation

Srinivasan & Bajorath, Cell Reports Physical Science 5, 102255
November 20, 2024
Published by Elsevier Inc.
https://doi.org/10.1016/j.xcrp.2024.102255

Cell Reports
Physical Science

CellPress
OPEN ACCESS

Article

Generation of dual-target compounds using a transformer chemical language model

Sanjana Srinivasan[1,2] and Jürgen Bajorath[1,2,3,4,*]

SUMMARY

Compounds with defined multi-target activity are candidates for the treatment of multi-factorial diseases. Such compounds are mostly discovered experimentally. Designing compounds with the desired activity against two targets is typically attempted by pharmacophore fusion. In addition, machine learning models can be derived for multi-target prediction of compounds or computational target profiling. Here, we introduce transformer-based chemical language model variants for the generative design of dual-target compounds. Alternative models were pre-trained by learning mappings of single- to dual-target compounds of increasing similarity. Different models were optimized for generating compounds with activity against pairs of functionally unrelated targets using a new cross-fine-tuning approach. Control models confirmed that pre-trained and fine-tuned models learned the chemical space of dual-target compounds. The final models were found to exactly reproduce known dual-target compounds excluded from model derivation. In addition, many structural analogs of such compounds were generated, further supporting the validity of the methodology.

INTRODUCTION

In drug discovery, target-specific compounds have been a primary focal point since the 1980s, when advances in molecular and structural biology triggered the introduction of the "one drug, one target" paradigm.[1] However, following the advent of systems biology in drug discovery,[2] evidence was also mounting that perturbed signaling and pharmacological networks were responsible for various multi-factorial diseases such as cancer[3,4] and that successful treatment of such diseases often depended on therapeutic intervention of multiple targets.[4,5] Such insights triggered the inception of the polypharmacology concept in drug discovery that was formally introduced in 2006,[5] encompassing the use of compounds with desirable multi-target (MT) activity and the exploration and exploitation of ensuing pharmacological effects.[5–7] Of note, MT engagement of drugs also is a major source of adverse side effects that require careful consideration in drug discovery.[7,8] While MT activity of some central nervous system drugs was already known since the late 1980s,[9,10] compounds with well-defined MT activity ("multi-specificity") became of central relevance for polypharmacology-oriented drug discovery.

Polypharmacological compounds were often discovered serendipitously, for instance in biological screening or target profiling campaigns, but were also of high interest in molecular design.[11,12] In computer-aided medicinal chemistry, the most popular design strategy for MT compounds (MT-CPDs) has thus far been combining or merging target-specific pharmacophores, a knowledge-based approach for which different computational protocols are available.[11–13] For example, pharmacophore fusion can

[1]Department of Life Science Informatics and Data Science, B-IT, University of Bonn, Friedrich-Hirzebruch-Allee 5/6, 53115 Bonn, Germany

[2]Lamarr Institute for Machine Learning and Artificial Intelligence, University of Bonn, Friedrich-Hirzebruch-Allee 5/6, 53115 Bonn, Germany

[3]Limes Institute - Program Unit Chemical Biology and Medicinal Chemistry, University of Bonn, Friedrich-Hirzebruch-Allee 5/6, 53115 Bonn, Germany

[4]Lead contact

*Correspondence: bajorath@bit.uni-bonn.de

https://doi.org/10.1016/j.xcrp.2024.102255

be attempted by designing hybrid compounds. Alternatively, multiple pharmacophore models might be used for computational compound screening concentrating on shared hits.[11–13] Other virtual screening protocols are also applicable. For example, machine learning (ML) models were used to search for compounds mimicking natural molecules with MT activity.[14] In addition, in computational target profiling, multiple independently derived single-target-based ML classifiers or multi-task (target) models can be used to search for hits with putative MT activity.[15,16] Target profiling is related to the ML-based prediction of drug-target interactions,[17–19] including proteochemometric modeling,[20] which also aids in the identification of MT-CPDs. Moreover, in structure-based design, cross-docking screens are applicable to prioritize compounds that can be accommodated in different binding sites based on high interaction and shape complementarity.[20,21]

While it is fair to state that pharmacophore-based approaches have thus far dominated MT-CPD design, ML offers new opportunities, supported by increasing numbers of compounds with multiple target annotations that are becoming available for learning. Notably, the most interesting MT-CPDs are not those that are active against closely related targets from the same family, which is often observed (due to a lack of compound specificity), but others that are active against structurally and functionally unrelated targets. For instance, a systematic analysis of public biological screening assays identified more than 1,000 compounds that were tested against at least 100 human targets and active against 10 or more targets from two or more different classes.[22] In addition, MT activities of nearly 2,000 drugs related to adverse side effects were systematically evaluated in 200 assays,[23] and the data have been made publicly available. Furthermore, a recent analysis of public compound data from medicinal chemistry identified nearly 700 compounds with activity against at least three targets from two different protein classes based on high-confidence target annotations.[24] Moreover, a systematic search in the ChEMBL compound database[25] uncovered 170 target pairs, involving 137 unique human targets, for which at least 100 dual-target compounds (DT-CPDs) and (50+50) corresponding single-target compounds (ST-CPDs) with high-confidence activity data were available.[26] Taken together, these findings illustrate that currently available MT-CPDs provide a substantial resource for follow-up analysis in medicinal chemistry and drug design. However, derivation and assessment of ML models to detect or generate new MT-CPDs is currently essentially restricted to compounds with activity against target pairs. For target triplets, numbers of available MT-CPDs are rapidly declining,[26] preventing meaningful ML model building in most cases. For the 170 target pairs referred to above, random forest classification models were built to distinguish DT- from corresponding ST-CPDs. Most of these models reached high prediction accuracy with a median of greater than 80%. Explainable ML analysis then revealed that these overall accurate pair-based predictions were determined by structural features that differed for target pairs and were present in DT- but absent in corresponding ST-CPDs.[26]

While classification models for MT-CPDs can be used for computational screening, deep generative ML models for drug discovery[27] can also be adapted for the design of compounds with desired MT activity. Thus, generative modeling further extends the capacity of ML for discovering MT-CPDs. However, compared to pharmacophore fusion or virtual screening, generative design of compounds for polypharmacology is still in its early stages, with only a few currently available studies. For instance, a general-purpose generative recurrent neural network (RNN), termed REINVENT 2.0,[28] was fine-tuned via transfer learning using 1,000 randomly selected MT-CPDs.[29,30] For this work, MT-CPDs were extracted from publicly available screening data.

Cell Reports
Physical Science
Article

CellPress
OPEN ACCESS

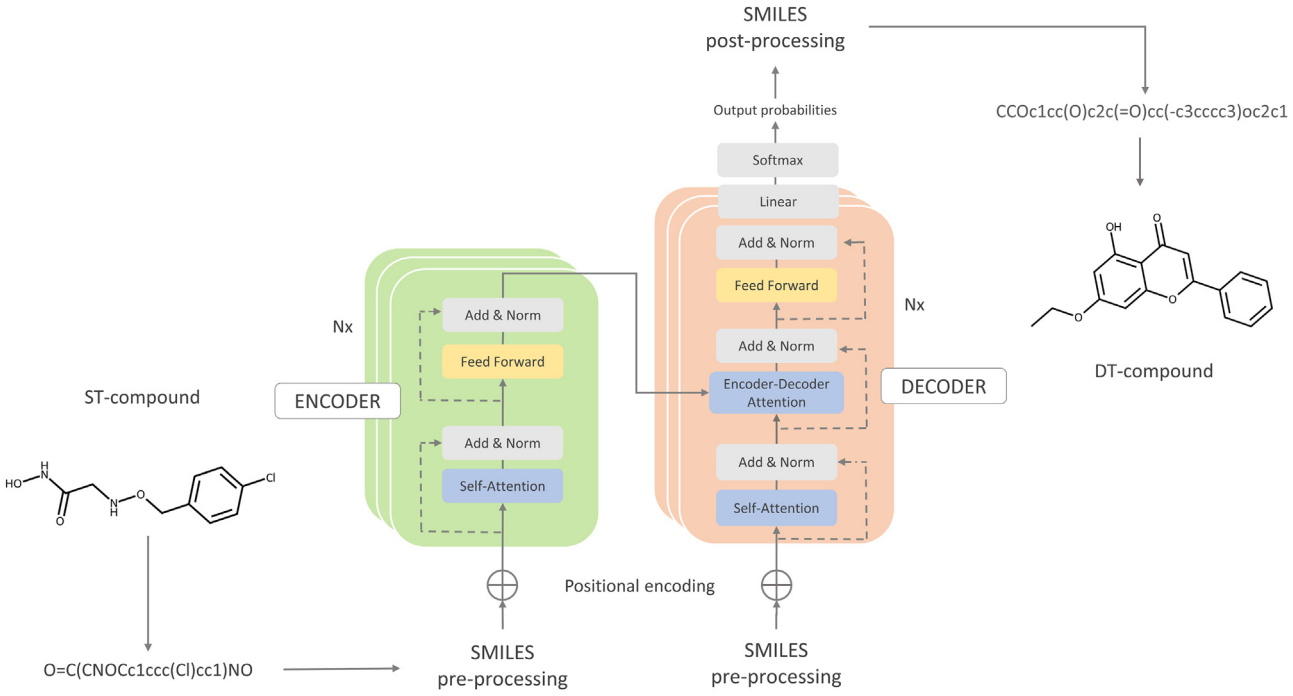

Figure 1. Transformer model
The architecture of the transformer-based CLM is schematically illustrated.

Specifically, MT-CPDs were selected that were active against five or more different targets. Of note, for this study, MT activity was generalized, without distinguishing between target combinations (otherwise the analysis would not have been feasible). Corresponding ST-CPDs were only active against one of these targets. During fine-tuning, the model produced increasing proportions of candidate compounds that were predicted to be MT-CPDs by an independently derived classifier, and the fine-tuned model was found to reproduce individual test set MT-CPDs and also generated structural analogs of such compounds.[29] Furthermore, with DrugEx v.2, another RNN was used in combination with transfer learning and multi-objective optimization to design DT-CPDs with a low likelihood of anti-target activity.[31]

In this study, we further expand predictive modeling of candidate compounds for MT engagement and introduce transformer models for the generation of DT-CPDs for pre-defined target pairs. Transformers[32] represent a preferred deep learning architecture for chemical language models (CLMs)[33,34] applied to learn molecular conversions. In this context, the design of DT-CPDs is perceived and executed as a machine translation task, as detailed in the following.

RESULTS AND DISCUSSION

Computational framework

A transformer-based CLM was derived to learn mappings of ST- to DT-CPDs (see experimental procedures for details). The CLM architecture is illustrated in Figure 1. For a target pair (A, B), ST-CPDs active against target A were mapped to DT-CPDs active against both A and B. During pre-training, ST-/DT-CPDs from 75,274 target pairs were used. For fine-tuning, six target pairs from distinct protein families not included in pre-training were selected, as reported in Table 1. The CLM was evaluated by determining its ability to exactly reproduce test DT-CPDs, representing the most stringent criterion for model validation. Three versions of the

Please cite this article in press as: Srinivasan and Bajorath, Generation of dual-target compounds using a transformer chemical language model, Cell Reports Physical Science (2024), https://doi.org/10.1016/j.xcrp.2024.102255

Table 1. Target pairs for fine-tuning

				DT-CPDs		
Target A ID	Target B ID	Target A	Target B	0%	25%	50%
220	2039	acetylcholinesterase (ACE)	monoamine oxidase B (MOB)	76	75	59
220	4822	ACE	beta-secretase 1 (BS1)	74	73	45
2971	1865	tyrosine-protein kinase JAK2 (JAK)	histone deacetylase 6 (HD6)	50	50	46
251	325	adenosine A2a receptor (A2R)	HD1	39	39	28
4630	220	serine/threonine-protein kinase CHK1 (CHK)	ACE	30	30	30
264	222	histamine H3 receptor (H3R)	norepinephrine transporter (NER)	27	27	27

For targets A and B of each pair, ChEMBL IDs and names are provided (abbreviations in parentheses are used in the following). In addition, the total number of DT-CPDs is reported for the 0%, 25%, and 50% similarity threshold datasets (see experimental procedures).

CLM were developed based on sets of ST- and DT-CPDs with different degrees of similarity (0%, 25%, and 50%). Furthermore, as a ST-CPD control (ST-CTRL), test sets of ST-/ST-CPD pairs were derived (see experimental procedures).

Pre-trained models

The three (0%, 25%, and 50%) versions of the pre-trained model were tested on hold-out sets comprising 10% of the unique DT-CPDs covering the entire target space. As reported in Table 2, the validity of the sampled strings across all test compounds was consistently close to 100%. For the 0% and 25% model versions, most of the sampled compounds were structurally unique, with mean uniqueness values of 98.4% and 82.6%, respectively. However, for the 50% model pre-trained on structurally very similar compounds, uniqueness was reduced to 27.1%, showing that this model often sampled multiple copies of the same compound. This was a likely consequence of restricting the chemical space for learning to very similar compounds, thereby also increasing the likelihood of re-sampling. On the other hand, the proportion of generated compounds that were structurally novel (not contained in training data) decreased from 35.2% (50% model) to 17.8% (25%) and 5.7% (0%), indicating that increasing the training compound similarity favored the generation of novel compounds, an unexpected finding at a first glance. However, this observation could be rationalized as follows: for very similar training compounds, confined structural modifications often generated different analogs formally qualifying as new compounds, which extended memorization of training compounds. Importantly, as also reported in Table 2, all three versions of the model accurately reproduced hundreds of test DT-CPDs not encountered during training, ranging from 31.8% (50% model) and 21.5% (25%) to 13.2% (0%). Hence, increasing the structural similarity of training and test compounds increased the probability of reproducing DT-CPDs from test ST-/DT-CPD pairs, as one might expect. The difference in reproducibility between the models might conceivably be attributed to the difference in the number of times an ST-CPD was used as input in pairs, thereby potentially increasing the probability of reproducing a DT-CPD. However, the number of times an ST-CPD

Table 2. Evaluation of pre-trained models

Model	ST-/DT-CPD pairs	DT-CPDs	Rep_DT-CPDs	R %	Similarity %	Validity %	Uniqueness %	Novelty %
0%	18,954	2,895	384	13.2	14.5	99.4	98.4	5.7
25%	18,249	2,416	521	21.5	27.9	98.0	82.6	17.8
50%	19,370	2,410	768	31.8	53.5	97.6	27.1	35.2

For each model version, the number of test ST-/DT-CPD pairs, unique DT-CPDs, accurately reproduced DT-CPDs (Rep_DT-CPDs), and the corresponding reproducibility (R) are reported. In addition, the average Tanimoto similarity of newly generated and all test DT-CPDs (similarity), the mean percentage of valid sampled SMILES strings (validity), unique sampled structures (uniqueness), and novel structures not contained in training data (novelty) are provided. Similarity, validity, uniqueness, and novelty values are averages over all test instances.

Please cite this article in press as: Srinivasan and Bajorath, Generation of dual-target compounds using a transformer chemical language model, Cell Reports Physical Science (2024), https://doi.org/10.1016/j.xcrp.2024.102255

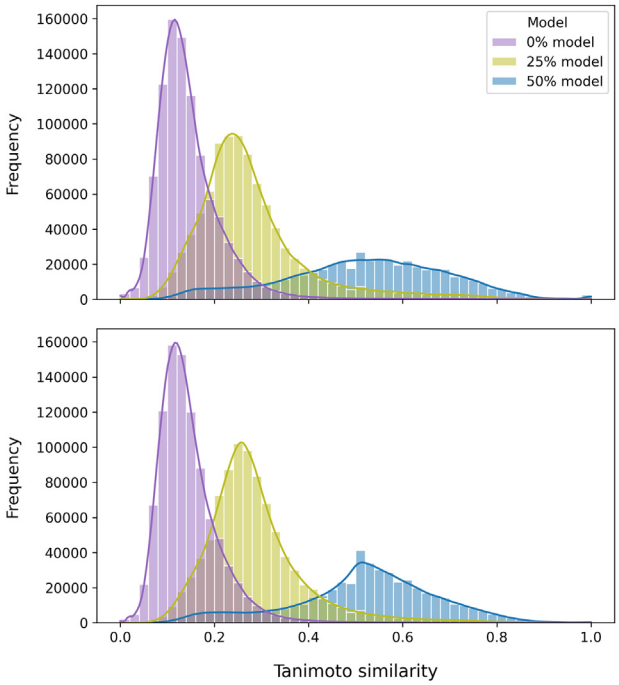

Figure 2. Similarity value distributions for pre-trained models

The distributions of ECFP4 Tanimoto similarity values from pairwise comparisons of candidate compounds generated with different model versions and corresponding test compounds are shown for DT-CPDs (top) and ST-CPDs (bottom).

was sampled was comparable for the 0%, 25%, and 50% models, with mean values of 1.3, 1.6, and 1.5, respectively (and a median value of 1 for all three models). Moreover, in the ST-CTRL, the corresponding reproducibility of test ST-CPDs was only low, i.e., 2.7% (50% model), 0.7% (25%), and 0.1% (0%). Thus, taken together, these findings provided proof of concept for the approach to generate known DT-CPDs from input ST-CPDs and demonstrated that the calculations were not determined by global compound similarity relationships but that the models indeed learned the chemical space of DT-CPDs.

Figure 2 shows the distributions of similarity values for pairwise comparisons of compounds generated with the different model versions and test ST- and DT-CPDs from all ST-/DT-CPD pairs. The distributions are similar for ST- and DT-CPDs. The modes of the distributions are close to the respective compound similarity threshold values but increasingly broaden with increasing similarity constraints. Hence, the 50% model also generated a substantial portion of compounds with lower similarity than test input and output compounds. For the model without compound similarity constraints, most compounds had close to 20% similarity to test compounds.

Figure 3 shows the distributions of synthetic accessibility (SA) scores calculated for all available DT-CPDs and compounds generated with the pre-trained 50% model. For known DT-CPDs and newly generated compounds, median SA values of 3.9 and 2.7 were obtained, respectively, indicating the slightly easier SA of model compounds.

Taken together, the findings in Figures 2 and 3 show that the pre-trained models produced chemically reasonable candidate compounds, providing a sound basis for fine-tuning.

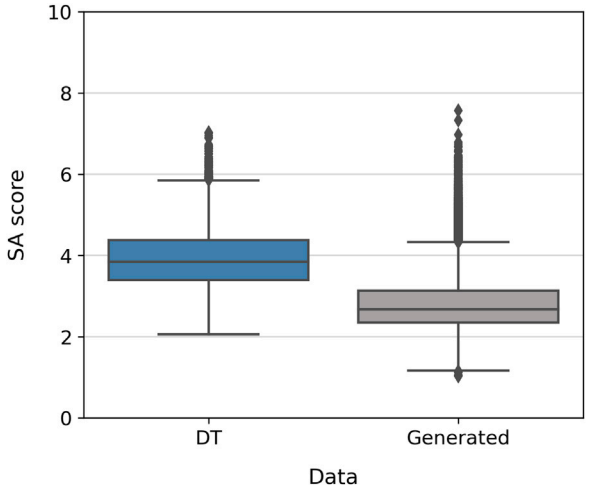

Figure 3. SA scores

Boxplots show the distributions of SA scores for DT-CPDs and candidate compounds generated with the 50% model. The boxes represent the upper and lower quartiles, the horizontal line the median value, and the whiskers the minimum and maximum values. Statistical outliers are shown as diamond symbols.

Fine-tuned models

For the pre-trained 0%, 25%, and 50% model versions, systematic cross-fine-tuning using 0%, 25%, and 50% fine-tuning datasets of the six pairs of unrelated targets in Table 1 was carried out. The results obtained for fine-tuning the pre-trained 0% model with the 0%, 25%, and 50% datasets are reported in Table 3.

Between 10 and 30 test DT-CPDs (in the 0% and 25% fine-tuning datasets) or between 10 and 23 (50% fine-tuning dataset) were available to test the fine-tuned model versions. After fine-tuning, the pre-trained model without a similarity constraint (0%) accurately reproduced one or more test DT-CPDs in 16 of 18 cases except for two target pairs when 0% fine-tuning data were used. For 0% fine-tuning data, one DT-CPD at most was reproduced for any target pair. These observations were made for all three pre-trained model versions. However, increasing similarity of the fine-tuning compounds increased the number (and proportion) of DT-CPDs reproduced with the pre-trained 0% model, with up to three and seven compounds for 25% and 50% fine-tuning data, respectively. Hence, for the confined numbers of DT-CPDs available for fine-tuning, the similarity of these compounds played a critically important role in focusing models on target pairs, as expected.

Table 4 reports the cross-fine-tuning results for the 25% and 50% pre-trained models. Given the low reproducibility observed for the 0%-0% model, 0% fine-tuning data were not considered for these models.

Fine-tuning differentiated between the pre-trained 0% and 25% models (Tables 3 and 4). The 25%-25% model clearly reproduced more test DT-CPDs than the 0%-25% model, with up to nine (compared to three) compounds and an average of 5.7 DT-CPDs per target pair, compared to 2.0 for the 0%-25% model. A similar observation was made for the 25%-50% model, which reproduced on average 5.3 test DT-CPDs per target pair, compared to 3.7 for the 0%-50% model. The numbers of reproduced test DT-CPDs were similar for the 25%-25% and 25%-50% models. However, given the smaller number of available test DT-CPDs in the 50% fine-tuning

Table 3. Fine-tuning of the 0% pre-trained model

	0% Fine-tuning			25% Fine-tuning			50% Fine-tuning		
Target pairs	DT-CPDs	Rep_DT-CPDs	R %	DT-CPDs	Rep_DT-CPDs	R %	DT-CPDs	Rep_DT-CPDs	R %
ACE-MOB	30	1	3.3	30	3	10.0	23	3	13.0
ACE-BS1	29	1	3.4	29	2	6.9	18	7	38.8
JAK-HD6	20	1	5.0	20	1	5.0	18	4	22.2
A2R-HD1	15	0	0.0	15	1	6.6	11	5	45.4
CHK-ACE	12	1	8.3	12	2	16.6	12	1	8.3
H3R-NER	10	0	0.0	10	3	30.0	10	2	20.0

For each target pair and fine-tuning dataset, the number of unique test DT-CPDs, the number of DT-CPDs exactly reproduced (Rep_DT-CPDs), and the reproducibility (R %) after fine-tuning are reported.

dataset, the mean reproducibility of the 25%-50% model (36.4%) was slightly higher than that of the 25%-25% model (30.6%). For the pre-trained 50% model, results after fine-tuning were very similar to the 25% model, as also reported in Table 4. The reproducibility of the different fine-tuned models across different trials is summarized and compared in Figure 4. Taken together, the results of cross-fine-tuning analysis revealed that similarity of pre-training and fine-tuning compounds was required for the derivation of models for the effective prediction of DT-CPDs for different target pairs and that the 25% and 50% models had comparable predictive ability.

Table 5 reports the best fine-tuned model for each target pair, including two 50%-50% and two 25%-50% models as well as one 50%-25% and one 25%-25% model. For these models, the number of unique sampled candidate compounds ranged from 102 to 701, with an average of 395. The models achieved high reproducibility of DT-CPDs of 23.3%–63.6%, depending on the target pair, with a mean of 42.6%. Furthermore, in addition to exactly reproduced test DT-CPDs, Table 5 also reports the proportion of structural analogs of test DT-CPDs among all sampled candidates, ranging from 1.7% to 44.2%, with a mean of 16.8%. Thus, structural analogs were frequently generated. Finally, the independently derived target-pair-based balanced random forest (BRF) classifiers predicted from 55.6% to 84.7% of sampled candidate compounds to be DT-CPDs, with a mean of 68.5%, thus indicating an enrichment of putative DT-CPDs among the generated candidates. While the reproducibility of known DT-CPDs excluded from the training phase represented the most rigorous criterion for predictive ability, these independent predictions provided supporting, albeit hypothetical, evidence.

The high reproducibility of fine-tuned models and frequent generation of structural analogs of test-DT compounds were encouraging, given that the chemical space of

Table 4. Fine-tuning of the 25% and 50% pre-trained model

	25% Pre-training						50% Pre-training					
	25% Fine-tuning			50% Fine-tuning			25% Fine-tuning			50% Fine-tuning		
Target pairs	DT-CPDs	Rep_DT-CPDs	R %	DT-CPDs	Rep_DT-CPDs	R %	DT-CPDs	Rep_DT-CPDs	R %	DT-CPDs	Rep_DT-CPDs	R %
ACE-MOB	30	7	23.3	23	4	17.4	30	4	13.3	23	5	21.7
ACE-BS1	29	6	20.6	18	9	50.0	29	8	27.6	18	10	55.5
JAK-HD6	20	7	35.0	18	7	38.9	20	8	40.0	18	7	38.9
A2R-HD1	15	9	60.0	11	7	63.6	15	9	60.0	11	4	36.4
CHK-ACE	12	3	25.0	12	1	8.3	12	3	25.0	12	4	33.3
H3R-NER	10	2	25.0	10	4	40.0	10	2	25.0	10	2	20.0

For each target pair and the 25% and 50% fine-tuning datasets, the number of unique test DT-CPDs, the number of DT-CPDs reproduced (Rep_DT-CPDs), and the reproducibility (R %) after fine-tuning are reported.

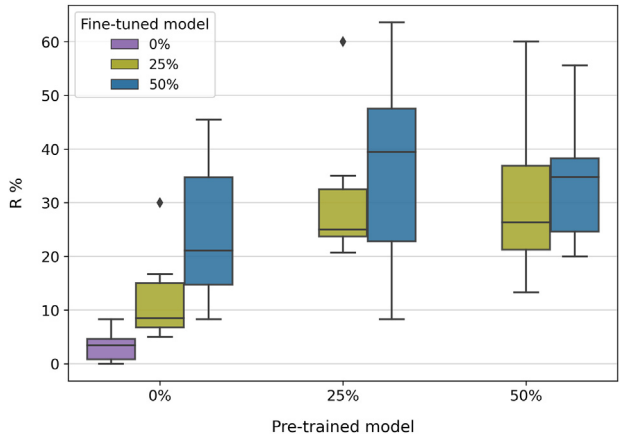

Figure 4. Reproducibility for fine-tuned models

Boxplots show the distribution of reproducibility (R) after fine-tuning for the different model versions.

DT-CPDs was sparsely populated and the likelihood to fully reproduce such compounds was intrinsically low. Moreover, for the 25%-25% and 50%-50% models, the ST-CTRL did not reproduce any test ST-CPDs for three and five of the six target pairs, respectively, and only a single or three ST-CPDs for the remaining pairs, thus clearly indicating that the fine-tuned models learned DT-CPDs.

Figure 5 shows exemplary predictions using 25%-25% models for different target pairs. In all cases, the test DT-CPD was correctly reproduced, and structural analogs of the DT-CPD were generated.

Limitations

Generative modeling of MT-CPDs is generally restricted by the limited availability of known MT-CPDs for learning, especially for unrelated targets, as discussed herein. Furthermore, detectable similarity between known ST- and DT-CPDs of a target pair is a pre-requisite for the derivation of generative models and effective fine-tuning, as revealed by our analysis. This also represents a general limitation, regardless of model architectures and prediction protocols. Learning of compound mappings implicitly capturing different properties must inevitably rely on detecting structural features and relationships giving rise to property changes.

Closing remarks

The prediction of compounds with defined MT activity is a topical issue in medicinal chemistry and drug design, given its immediate relevance for polypharmacology-oriented

Table 5. Preferred fine-tuned models for target pairs

Target pairs	Best model	R %	Analogs %	Predicted DT-CPDs %
ACE-MOB	25%-25%	23.3	11.5	61.7
ACE-BS1	50%-50%	55.5	1.7	62.8
JAK-HD6	50%-25%	40.0	16.8	84.7
A2R-HD1	25%-50%	63.6	20.6	76.4
CHK-ACE	50%-50%	33.3	44.2	70.0
H3R-NER	25%-50%	40.0	6.1	55.6

For each target pair, the best-performing fine-tuned model version is reported together with its reproducibility (R). In addition, for each model, "analogs" reports the proportion of structural analogs of test DT-CPDs among sampled candidate compounds and "predicted DT-CPDs" the proportion of sampled candidate compounds that were predicted to be DT-CPDs using target-pair-based BRF classifiers.

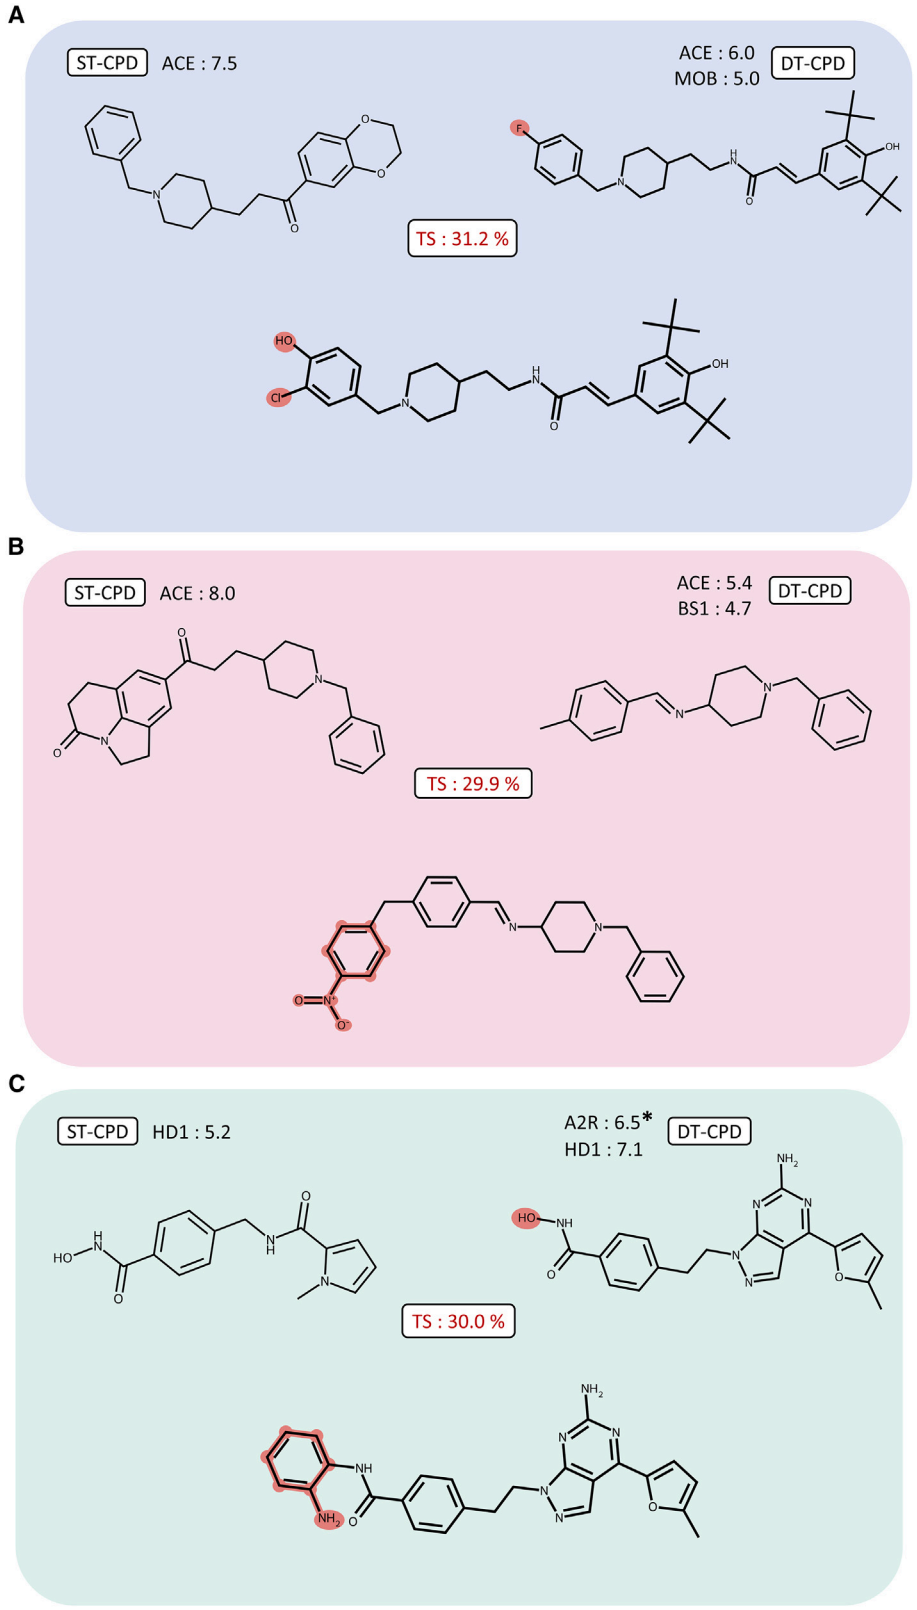

Figure 5. Compound generation

Shown are results of exemplary predictions using 25%-25% models for different target pairs. In each case, a test DT-CPD was exactly reproduced for an input ST-CPD, and structural analogs of the DT-CPD were generated. The input ST-CPD, output DT-CPD, and an exemplary structural analog (bottom) are shown for the (A) ACE-MOB, (B) ACE-BS1, and (C) A2R-HD1 target pairs. For the input ST-CPDs and output DT-CPDs, experimental negative-decadic logarithmic potency values for their targets (IC50 values except one K_d value indicated with an asterisk) and Tanimoto similarity (TS) are reported. In the DT-CPDs and the sampled structural analogs (bottom), the distinguishing substituents are highlighted. All structural analogs were novel compounds not encountered during training.

drug discovery. Compounds are often active against closely related targets (e.g., targets from the same family), but such compounds are not necessarily prime candidates for polypharmacology, except in cases where particular families are targeted (such as protein kinase families implicated in cancer). Rather, compounds with the potential to inhibit or antagonize targets with distinct functions are often desirable, for instance, to simultaneously interfere with signaling and metabolic pathways implicated in disease. For all practical purposes, the design of MT-CPDs is primarily focused on compounds with activity against two or, at most, three distantly related or unrelated targets. While pharmacophore fusion methods have played a dominant role in practical MT-CPD design thus far, advanced ML approaches are also considered but are still in their infancy for this design task. In this work, we have introduced transformer-based CLM variants for the generative design of DT-CPDs. Of note, for compounds with triple-target activity, there currently are too few available for systematic model derivation and assessment, as reported herein. Only for related protein kinases were a limited number of datasets with inhibitors having triple-kinase activity and corresponding single-kinase activity assembled,[35] given the wealth of compound activity data available for this target class. Our current models were pre-trained to cover a large target space and learn mappings of ST- to DT-CPDs at different similarity thresholds. The pre-trained models generated almost exclusively valid simplified molecular input line entry system (SMILES) strings and preferentially sampled unique compounds with favorable SA scores. The pre-trained model versions were then focused on DT-CPDs for pairs of functionally unrelated targets (excluded from pre-training) via cross fine-tuning applying compound similarity thresholds corresponding to those used during pre-training. A comparison of 0% and 25% pre-trained and fine-tuned models showed that similarity of DT-CPDs and corresponding ST-CPDs was required for the derivation of predictive models, as one would expect. In the absence of detectable similarity relationships, learning compound mappings could hardly lead to the generation of candidate compounds with shared features. Importantly, ST-CTRL calculations confirmed that the pre-trained and fine-tuned models successfully learned mappings of ST- to DT-CPDs as the basis for the generative design of new candidates. The final fine-tuned models achieved, in part, unexpectedly high reproducibility of known DT-CPDs not encountered during training, representing the most stringent test for a generative model. In addition, the fine-tuned models produced varying amounts of structural analogs of known DT-CPDs (and large numbers of candidates predicted to be DT-CPDs by an independent classifier), thus further supporting their predictive ability. Taken together, our findings indicate that CLM variants for DT-CPD design presented herein are capable of generating candidate compounds that should merit careful consideration in drug discovery. To these ends, the pre-trained model versions, such as the 25% model, are immediately applicable for fine-tuning with compounds for target pairs of interest.

EXPERIMENTAL PROCEDURES

Compounds and target-pair-based datasets

High-confidence compound activity data were extracted from ChEMBL[25] (release 33). Compounds with a molecular mass of less than 1,000 Da and numerically specified (IC_{50}, K_d, or K_i) potency values ranging from 10 μM to 10 pM were selected. Compound-target interactions with multiple potency values were only retained if they fell within one order of magnitude. Direct interactions (target relationship

Please cite this article in press as: Srinivasan and Bajorath, Generation of dual-target compounds using a transformer chemical language model, Cell Reports Physical Science (2024), https://doi.org/10.1016/j.xcrp.2024.102255

type "D") with single human target proteins and an assay confidence score of 9 were also required. Compounds flagged with comments such as "inactive," "not active," "inconclusive," "potential transcription error," and "potential author error" were discarded. In addition, public filters were applied to exclude potential false positive activity annotations due to chemical liability, assay interference, or colloidal aggregator effects.[25,36–40] Finally, compounds with undesirable target annotations such as cytochrome 450 isoforms, hERG, or albumin were omitted.

Data curation resulted in the selection of 120,195 unique qualifying compounds with activity against 1,747 unique target proteins that were assigned to 329 families based on the UniProt classification.[41]

For all possible target pairs, a search for DT-CPDs was carried out. DT-CPDs were detected for 75,280 target pairs involving 1,457 unique targets from 231 UniProt families. A confined subset of 7,747 target pairs was formed by targets from different families. For each qualifying target pair, subsets of ST-/DT-CPD pairs with at least 50% or 25% Tanimoto similarity[42] were generated based on the extended connectivity fingerprint with bond diameter 4 (ECFP4)[43] using its 2,048-bit version generated using RDKit.[36]

For pre-training, ST-/DT-CPD pairs from all target pairs were pooled. Compound pairs in the pre-training datasets contained 42,328 unique ST- and 24,109 unique DT-CPDs.

Generative model

Data mapping and tokenization

The key task for our transformer-based CLM was to learn target-pair-based mappings of ST- to DT-CPDs that were represented as SMILES strings[44] (generated using RDKit). Thus, for a given target A in a pair (A, B), mapping of ST-CPDs active against A to DT-CPDs with activity against targets A and B was learned:

$$(ST_A - CPD) \rightarrow (DT_{[A,B]} - CPD).$$

Tokenization of SMILES strings[45] yielded a vocabulary consisting of 38 individual tokens, primarily composed of single-character tokens and also two-character tokens like "Br," and "'Cl," as well as bracketed tokens like "[nH]" and "[O-]." Additionally, "start" and "end" tokens were used to represent the start and end of strings, respectively.

Model architecture and implementation

A transformer encoder-decoder architecture[32] was adopted from a previously reported CLM derived to predict activity cliffs.[45] The architecture is illustrated in Figure 1. The encoder stack consists of six identical layers, each composed of two sub-layers, including a multi-head self-attention sub-layer and a fully connected feedforward neural network (FFN) sub-layer. The multi-head self-attention mechanism facilitates the retention of long-term memory by applying attention functions in parallel over different segments of the input sequence. Subsequently, the encoder processes the tokenized input data and encodes it into a continuous representation, which is then utilized by the decoder for generating the output sequence. The decoder stack also consists of six identical layers. In addition to the two sub-layers analogous to those in the encoder, the decoder constitutes a third multi-head attention sub-layer responsible for processing the output from the encoder (termed encoder-decoder attention). The self-attention sub-layer of the decoder also

incorporates masking to prevent the model from attending to the information from previous positions.[32]

Model variants were implemented using PyTorch[46] with default parameter settings including a dropout rate of 0.1, a batch size of 64, and a learning rate of 0.0001. The model was configured with a dimensionality of 256 and featured eight attention heads per layer, with label smoothing set to 0.

The training process included pre-training and fine-tuning. During training, the model was subjected to at least 100 epochs using the Adam optimizer[47] with its default settings. The Kullback-Leibler (KL)-divergence loss function[48] was employed to guide the optimization process. The model's state was saved after each epoch, and the version with the lowest validation loss was selected as the final model. In the decoder, the softmax function was used to generate probability distributions for tokens. Output compounds were generated through multinomial sampling, and the resulting tokens were re-converted to SMILES strings. For each input ST-CPD, a maximum of 50 unique and valid SMILES strings were sampled over a maximum of 100 trials.

For sampled output compounds, SA scores[49] were calculated and compared to the training and test set compounds. The SA scores ranged from 1 to 10, indicating decreasing SA.

Pre-training

For each of the 75,274 target pairs (excluding six pairs for fine-tuning), the 50% similarity subset was selected first. Then, the same number of ST-/DT-CPD pairs with 25% similarity was randomly sampled. In addition, an equally sized subset with no similarity constraint was generated by random sampling. These subsets of constant size represented the 50%, 25%, and "no similarity" (0%) datasets for pre-training (using 90% of the DT-CPDs) and initial testing (10%) of three different versions of the model. For testing, the ST-CPD of each available ST-/DT-CPD pair was used as input, and output compounds were sampled as specified above. The large number of target pairs used for pre-training ensured that the model comprehensively charted chemical space defined by DT-CPDs.

Furthermore, as ST-CTRL, test sets exclusively consisting of ST-CPDs were generated. ST-CPDs active against target A of a pair (A, B) were combined with randomly selected ST-CPDs with activity against any other target except A or B. Therefore, 50%, 25%, and 0% datasets of ST-/ST-CPD pairs of the same size as the corresponding ST-/DT-CPD datasets were generated and used to test the three pre-trained versions of the model. This control was implemented to verify that the ST-/DT-CPD models learned the DT-CPD chemical space rather than merely generating candidate compounds meeting the similarity criteria.

Fine-tuning

For fine-tuning, six target pairs from different families were selected, as reported in Table 1, and excluded from pre-training. These target pairs were selected because the targets belonged to functionally unrelated protein families and shared a sufficient number of DT-CPDs for model derivation and evaluation, which were exceptions rather than the rule. For pairs of unrelated targets, which are most interesting from a polypharmacology perspective, typically only small numbers of DT-CPDs (if at all) were available for learning and testing, which limited the ability to derive generative models for such target pairs. For fine-tuning, all

available DT-CPDs were included in the 50%, 25%, and 0% datasets, as reported in Table 1. Cross fine-tuning was systematically carried out, during which each pre-trained model was fine-tuned separately with datasets at different similarity levels. Hence, cross fine-tuning yielded nine models with varying (pre-training)-(fine-tuning) data combinations (i.e., 0%-0%, 0%-25%, 0%-50%, 25%-0%, 25%-25%, 25%-50%, 50%-0%, 50%-25%, and 50%-50%). For each target pair, 60% and 40% of available DT-CPDs were used for fine-tuning and testing, respectively. The pre-trained versions of the models were individually fine-tuned for each target pair.

For each target pair, an ST-CTRL test set was also generated in which all test ST-CPDs active against target A of a pair (A, B) were paired with randomly selected ST-CPDs, satisfying the respective similarity constraints and having activity against any target except A or B.

Model evaluation

For each input ST-CPD, the proportion of valid SMILES strings among all sampled strings was calculated and averaged. In addition, the proportions of structurally unique compounds that were sampled and novel compounds not contained in training data were determined.

As the most stringent criteria for evaluating predictions, the ability of a model *to exactly reproduce test DT-CPDs* was determined by calculating the reproducibility for test sets, defined as

$$\text{reproducibility (R)} = \frac{\text{number of reproduced DT} - \text{CPDs}}{\text{number of test DT} - \text{CPDs}} \times 100.$$

To complement the reproducibility measure, we also identified structural analogs of test DT-CPDs among newly generated candidates. Structural analogs represent compounds that share the same core structure and are only distinguished by one or more varying substituents (R-groups). Structural analogs of test compounds were systematically identified using the compound-core-relationship (CCR) algorithm.[50]

In addition, for all sampled candidate compounds, pairwise Tanimoto similarity was calculated and compared to test sets.

BRF classifiers

To further evaluate sampled candidate compounds, BRF classification models were developed for the six fine-tuning target pairs using scikit-learn[51] and imbalanced-learn[52] to distinguish between DT- and corresponding ST-CPDs. For each target pair, the BRF model was trained using 75% of the ST- and DT-CPDs. The classification performance of these models was evaluated using a balanced test set comprising the remaining 25% of the data. The balanced accuracy[53] (BA) scores for all models were consistently above 0.72, with a mean BA of 0.85.

RESOURCE AVAILABILITY

Lead contact

Further information and requests should be directed to the lead contact, Jürgen Bajorath (bajorath@bit.uni-bonn.de).

Materials availability

All compounds and associated data were obtained from ChEMBL (release 33) and UniProt (release 2024_01).

Data and code availability

All data and code used for our analysis can be accessed via the following link: https://uni-bonn.sciebo.de/s/y0macayDXoqxTVn.

ACKNOWLEDGMENTS

The authors thank Hengwei Chen and Martin Vogt for helpful discussions.

AUTHOR CONTRIBUTIONS

Conceptualization, J.B.; methodology, S.S. and J.B.; investigation, S.S. and J.B.; writing – original draft, S.S. and J.B.; writing – review & editing, S.S. and J.B.; supervision, J.B.

DECLARATION OF INTERESTS

The authors declare no competing interests.

Received: July 25, 2024
Revised: September 18, 2024
Accepted: September 26, 2024
Published: October 23, 2024

REFERENCES

1. Morphy, R., Kay, C., and Rankovic, Z. (2004). From magic bullets to designed multiple ligands. Drug Discov. Today *9*, 641–651. https://doi.org/10.1016/S1359-6446(04)03163-0.

2. Butcher, E.C., Berg, E.L., and Kunkel, E.J. (2004). Systems biology in drug discovery. Nat. Biotechnol. *22*, 1253–1259. https://doi.org/10.1038/nbt1017.

3. Hopkins, A.L. (2007). Network pharmacology. Nat. Biotechnol. *25*, 1110–1111. https://doi.org/10.1038/nbt1007-1110.

4. Ainsworth, C. (2011). Networking for new drugs. Nat. Med. *17*, 1166–1168. https://doi.org/10.1038/nm1011-1166.

5. Hopkins, A.L., Mason, J.S., and Overington, J.P. (2006). Can we rationally design promiscuous drugs? Curr. Opin. Struct. Biol. *16*, 127–136. https://doi.org/10.1016/j.sbi.2006.01.013.

6. Anighoro, A., Bajorath, J., and Rastelli, G. (2014). Polypharmacology: Challenges and Opportunities in Drug Discovery: Miniperspective. J. Med. Chem. *57*, 7874–7887. https://doi.org/10.1021/jm5006463.

7. Peters, J.-U. (2013). Polypharmacology – Foe or Friend? J. Med. Chem. *56*, 8955–8971. https://doi.org/10.1021/jm400856t.

8. Berger, S.I., and Iyengar, R. (2011). Role of systems pharmacology in understanding drug adverse events. WIREs Mechanisms of Disease *3*, 129–135. https://doi.org/10.1002/wsbm.114.

9. Benek, O., Korabecny, J., and Soukup, O. (2020). A Perspective on Multi-target Drugs for Alzheimer's Disease. Trends Pharmacol. Sci. *41*, 434–445. https://doi.org/10.1016/j.tips.2020.04.008.

10. Weston-Green, K. (2022). Antipsychotic Drug Development: From Historical Evidence to Fresh Perspectives. Front. Psychiatry *13*, 903156. https://doi.org/10.3389/fpsyt.2022.903156.

11. Proschak, E., Stark, H., and Merk, D. (2019). Polypharmacology by Design: A Medicinal Chemist's Perspective on Multitargeting Compounds. J. Med. Chem. *62*, 420–444. https://doi.org/10.1021/acs.jmedchem.8b00760.

12. Li, X., Li, X., Liu, F., Li, S., and Shi, D. (2021). Rational Multitargeted Drug Design Strategy from the Perspective of a Medicinal Chemist. J. Med. Chem. *64*, 10581–10605. https://doi.org/10.1021/acs.jmedchem.1c00683.

13. Moser, D., Wisniewska, J.M., Hahn, S., Achenbach, J., Buscató, E.l., Klingler, F.-M., Hofmann, B., Steinhilber, D., and Proschak, E. (2012). Dual-Target Virtual Screening by Pharmacophore Elucidation and Molecular Shape Filtering. ACS Med. Chem. Lett. *3*, 155–158. https://doi.org/10.1021/ml200286e.

14. Vamathevan, J., Clark, D., Czodrowski, P., Dunham, I., Ferran, E., Lee, G., Li, B., Madabhushi, A., Shah, P., Spitzer, M., and Zhao, S. (2019). Applications of machine learning in drug discovery and development. Nat. Rev. Drug Discov. *18*, 463–477. https://doi.org/10.1038/s41573-019-0024-5.

15. D'Souza, S., Prema, K.V., and Balaji, S. (2020). Machine learning models for drug–target interactions: current knowledge and future directions. Drug Discov. Today *25*, 748–756. https://doi.org/10.1016/j.drudis.2020.03.003.

16. Chen, H., Cheng, F., and Li, J. (2020). iDrug: Integration of drug repositioning and drug-target prediction via cross-network embedding. PLoS Comput. Biol. *16*, e1008040. https://doi.org/10.1371/journal.pcbi.1008040.

17. Bai, P., Miljković, F., John, B., and Lu, H. (2023). Interpretable bilinear attention network with domain adaptation improves drug–target prediction. Nat. Mach. Intell. *5*, 126–136. https://doi.org/10.1038/s42256-022-00605-1.

18. Bongers, B.J., IJzerman, A.P., and Van Westen, G.J.P. (2019). Proteochemometrics – recent developments in bioactivity and selectivity modeling. Drug Discov. Today Technol. *32–33*, 89–98. https://doi.org/10.1016/j.ddtec.2020.08.003.

19. Grisoni, F., Merk, D., Friedrich, L., and Schneider, G. (2019). Design of Natural-Product-Inspired Multitarget Ligands by Machine Learning. ChemMedChem *14*, 1129–1134. https://doi.org/10.1002/cmdc.201900097.

20. Rastelli, G., and Pinzi, L. (2015). Computational polypharmacology comes of age. Front. Pharmacol. *6*, 157. https://doi.org/10.3389/fphar.2015.00157.

21. Caballero, J. (2021). The latest automated docking technologies for novel drug discovery. Expert Opin. Drug Discov. *16*, 625–645. https://doi.org/10.1080/17460441.2021.1858793.

22. Feldmann, C., Miljković, F., Yonchev, D., and Bajorath, J. (2019). Identifying Promiscuous Compounds with Activity against Different Target Classes. Molecules *24*, 4185. https://doi.org/10.3390/molecules24224185.

23. Sutherland, J.J., Yonchev, D., Fekete, A., and Urban, L. (2023). A preclinical secondary pharmacology resource illuminates target-adverse drug reaction associations of marketed drugs. Nat. Commun. *14*, 4323. https://doi.org/10.1038/s41467-023-40064-9.

24. Srinivasan, S., and Bajorath, J. (2024). Systematic identification and characterization of compounds with reliable activity against multiple target proteins from different classes. Results Chem. *7*, 101376. https://doi.org/10.1016/j.rechem.2024.101376.

25. Gaulton, A., Hersey, A., Nowotka, M., Bento, A.P., Chambers, J., Mendez, D., Mutowo, P., Atkinson, F., Bellis, L.J., Cibrián-Uhalte, E., et al. (2017). The ChEMBL database in 2017. Nucl. Acids Res. *45*, D945–D954. https://doi.org/10.1093/nar/gkw1074.

26. Feldmann, C., Philipps, M., and Bajorath, J. (2021). Explainable machine learning

predictions of dual-target compounds reveal characteristic structural features. Sci. Rep. *11*, 21594. https://doi.org/10.1038/s41598-021-01099-4.

27. Tong, X., Liu, X., Tan, X., Li, X., Jiang, J., Xiong, Z., Xu, T., Jiang, H., Qiao, N., and Zheng, M. (2021). Generative Models for De Novo Drug Design. J. Med. Chem. *64*, 14011–14027. https://doi.org/10.1021/acs.jmedchem.1c00927.

28. Blaschke, T., Arús-Pous, J., Chen, H., Margreitter, C., Tyrchan, C., Engkvist, O., Papadopoulos, K., and Patronov, A. (2020). REINVENT 2.0: An AI Tool for De Novo Drug Design. J. Chem. Inf. Model. *60*, 5918–5922. https://doi.org/10.1021/acs.jcim.0c00915.

29. Blaschke, T., and Bajorath, J. (2022). Fine-tuning of a generative neural network for designing multi-target compounds. J. Comput. Aided Mol. Des. *36*, 363–371. https://doi.org/10.1007/s10822-021-00392-8.

30. Blaschke, T., and Bajorath, J. (2021). Compound Dataset and Custom Code for Deep Generative multi-target Compound Design. Future Sci. OA *7*, FSO715. https://doi.org/10.2144/fsoa-2021-0033.

31. Liu, X., Ye, K., Van Vlijmen, H.W.T., Emmerich, M.T.M., IJzerman, A.P., and Van Westen, G.J.P. (2021). DrugEx v2: de novo design of drug molecules by Pareto-based multi-objective reinforcement learning in polypharmacology. J. Cheminform. *13*, 85. https://doi.org/10.1186/s13321-021-00561-9.

32. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., and Polosukhin, I. (2017). Attention Is All You Need. Preprint at arXiv. https://doi.org/10.48550/ARXIV.1706.03762.

33. Grisoni, F. (2023). Chemical language models for de novo drug design: Challenges and opportunities. Curr. Opin. Struct. Biol. *79*, 102527. https://doi.org/10.1016/j.sbi.2023.102527.

34. Yoshimori, A., Chen, H., and Bajorath, J. (2023). Chemical Language Models for Applications in Medicinal Chemistry. Future Med. Chem. *15*, 119–121. https://doi.org/10.4155/fmc-2022-0315.

35. Feldmann, C., and Bajorath, J. (2022). Differentiating Inhibitors of Closely Related Protein Kinases with Single- or Multi-Target Activity via Explainable Machine Learning and Feature Analysis. Biomolecules *12*, 557. https://doi.org/10.3390/biom12040557.

36. RDKit. (2024). RDKit: Open-Source Cheminformatics Software (GitHub). https://www.rdkit.org/

37. Irwin, J.J., Tang, K.G., Young, J., Dandarchuluun, C., Wong, B.R., Khurelbaatar, M., Moroz, Y.S., Mayfield, J., and Sayle, R.A. (2020). ZINC20—A Free Ultralarge-Scale Chemical Database for Ligand Discovery. J. Chem. Inf. Model. *60*, 6065–6073. https://doi.org/10.1021/acs.jcim.0c00675.

38. Baell, J.B., and Holloway, G.A. (2010). New Substructure Filters for Removal of Pan Assay Interference Compounds (PAINS) from Screening Libraries and for Their Exclusion in Bioassays. J. Med. Chem. *53*, 2719–2740. https://doi.org/10.1021/jm901137j.

39. Irwin, J.J., Duan, D., Torosyan, H., Doak, A.K., Ziebart, K.T., Sterling, T., Tumanian, G., and Shoichet, B.K. (2015). An Aggregation Advisor for Ligand Discovery. J. Med. Chem. *58*, 7076–7087. https://doi.org/10.1021/acs.jmedchem.5b01105.

40. Bruns, R.F., and Watson, I.A. (2012). Rules for Identifying Potentially Reactive or Promiscuous Compounds. J. Med. Chem. *55*, 9763–9772. https://doi.org/10.1021/jm301008n.

41. The, U.P.C., Bateman, A., Martin, M.-J., Orchard, S., Magrane, M., Ahmad, S., Alpi, E., Bowler-Barnett, E.H., Britto, R., Bye-A-Jee, H., et al. (2023). UniProt: the Universal Protein Knowledgebase in 2023. Nucl. Acids Res. *51*, D523–D531. https://doi.org/10.1093/nar/gkac1052.

42. Bajusz, D., Rácz, A., and Héberger, K. (2015). Why is Tanimoto index an appropriate choice for fingerprint-based similarity calculations? J. Cheminform. *7*, 20. https://doi.org/10.1186/s13321-015-0069-3.

43. Rogers, D., and Hahn, M. (2010). Extended-Connectivity Fingerprints. J. Chem. Inf. Model. *50*, 742–754. https://doi.org/10.1021/ci100050t.

44. Weininger, D. (1988). SMILES, a chemical language and information system. 1. Introduction to methodology and encoding rules. J. Chem. Inf. Comput. Sci. *28*, 31–36. https://doi.org/10.1021/ci00057a005.

45. Chen, H., Vogt, M., and Bajorath, J. (2022). DeepAC – conditional transformer-based chemical language model for the prediction of activity cliffs formed by bioactive compounds. Digit. Discov. *1*, 898–909. https://doi.org/10.1039/D2DD00077F.

46. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al. (2019). PyTorch: an imperative style, high-performance deep learning library. Adv. Neural Inf. Proc. Syst. *32*, 8026–8037. https://doi.org/10.48550/ARXIV.1912.01703.

47. Kingma, D.P., and Ba, J. (2014). Adam: a method for stochastic optimization. Preprint at arXiv. https://doi.org/10.48550/ARXIV.1412.6980.

48. Kullback, S., and Leibler, R.A. (1951). On information and sufficiency. Ann. Math. Statist. *22*, 79–86. https://doi.org/10.1214/aoms/1177729694.

49. Ertl, P., and Schuffenhauer, A. (2009). Estimation of synthetic accessibility score of drug-like molecules based on molecular complexity and fragment contributions. J. Cheminform. *1*, 8. https://doi.org/10.1186/1758-2946-1-8.

50. Naveja, J.J., Vogt, M., Stumpfe, D., Medina-Franco, J.L., and Bajorath, J. (2019). Systematic extraction of analogue series from large compound collections using a new computational compound–core relationship method. ACS Omega *4*, 1027–1032. https://doi.org/10.1021/acsomega.8b03390.

51. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., et al. (2011). Scikit-learn: machine learning in Python. J. Mach. Learn. Res. *12*, 2825–2830.

52. Lemaitre, G., Nogueira, F., and Aridas, C.K. (2016). Imbalanced-learn: a Python toolbox to tackle the curse of imbalanced datasets in machine learning. Preprint at arXiv. https://doi.org/10.48550/ARXIV.1609.06570.

53. Wei, Q., and Dunbrack, R.L. (2013). The role of balanced training and testing data sets for binary classifiers in bioinformatics. PLoS One *8*, e67863. https://doi.org/10.1371/journal.pone.0067863.